# Creating Energy-Efficient Data Centers

## ENERGY COUNTS

**BY JACK POUCHET**

Data centers and their power hungry servers have drawn the attention of the U.S. House of Representatives, which has called for a study on data center energy consumption. Server and chip manufacturers are addressing efficiency concerns with new products; however, rather than allowing technology-dependent organizations to do more with less, these advances are enabling them to do more with more.

As a result, the data center is moving from a cool controlled environment with power densities of 50 to 75 watts per square foot to large Easy Bake oven-like power densities of 200 to 300 watts per square foot and more. In fact, watts per square foot has become almost meaningless as a measure of data center capacity as densely packed racks make "white space" ineffective in handling cooling loads.

Of course, this is all happening while energy costs are rising, creating increased strain on operating budgets. Fortunately, there are ample opportunities to regain control of the data center and optimize availability and efficiency. In many cases, an immediate reduction in operating expense of 20 to 30% or higher can be achieved, depending upon current utilization rate, building design, and management practices.

> There are ample opportunities to regain control of the data center and optimize availability and efficiency.

### A Measure of Efficiency

Data center efficiency has many variables; however, simple metrics can be used to gauge operational performance against industry averages and organizational objectives. The measure put forth by Christian Belady and Christopher Malone, PhD, of HP, compares the power required to support data center operations (HVAC, chillers, UPS, lighting, and IT hardware) to the actual work performed by the IT hardware (the IT load) to determine a power utilization efficiency (PUE) ratio or, inversely, data center efficiency (DCE).

A "good" PUE is 1.6, (1.6 watts of electricity required for every 1 watt of computing) with a reasonable target of 2.0 for facilities just coming to grips with this concept. Recent studies funded by the California Energy Commission (CEC), (Energy Efficient Data Centers, December 2003 and Self Benchmarking Guide for Data Center Energy Performance version 1.0, May 2006) by Lawrence Berkeley National Labs indicate that data center PUEs vary widely. You can determine a PUE by comparing utility bills if the data center is separately metered with the IT hardware load. If the PUE is more than 2.0, it is likely there is opportunity to make relatively simple changes that can significantly reduce operating costs.

### Assessing the Problem

The major server OEMs, power, and cooling experts such as Emerson Network Power, and local engineering firms all provide walk-through data center assessments that can help identify opportunities to improve efficiency and performance. Costs and scope of these assessments vary widely, but, at minimum, they should include an examination of:

- Critical power systems, including generator(s), switchgear, utility entrance, fuel storage, UPS system and batteries, and PDUs
- HVAC systems including discharge and return air temperature/humidity, data log, filters, humidification/de-humidification cycles, zone/spot /supplemental cooling and chiller plant
- Room airflow, including tile placement, obstructions, rack arrangement, and aisle temperatures
- IT equipment loads and utilization
- Service records

The assessment should produce a fairly detailed report that can be used to justify capital expenses (if necessary), as well as identifying facility and process improvements. In many cases, assessments can identify simple and low-cost improvements such as:

- Improper floor tile layout (perforated tiles in the hot aisle)
- Blocked duct work, open duct work, blocked filters, large cable entry cut-outs
- Open space within racks that permit recirculation
- Under floor and overhead obstructions
- Mismatched HVAC to server airflows or unbalanced HVAC loading
- Competing HVACs (one or more in humidification mode with others in de-humidification)
- Unloaded or overloaded UPS units or PDUs.

With this information, a company can start to build an efficiency improvement plan. Work begins by evaluating the return-on-investment (ROI) for each recommendation in the audit based on the cost per kW/Mw of power. When calculating ROI, it may be wise to multiply the savings promised by the audit by 0.8 to account for variations in results across different facilities.

Then the list is ready to be prioritized based upon ease-of-implementation, total savings, or payback period. Depending upon where the data center is located, additional resources may

**HOT AISLE/ COLD AISLE APPROACH**

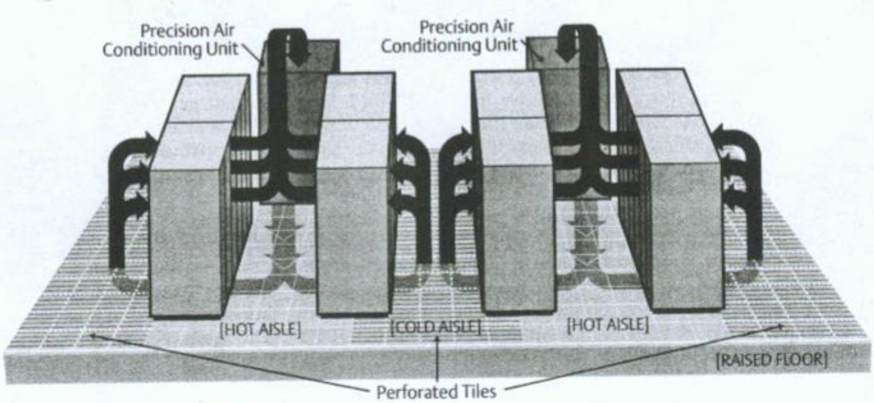




The hot aisle/cold aisle approach is effective in cooling data centers, but increasing power densities challenge the limits of current technologies and best practices.

| PS Efficiency (%) | Energy (1000, 1,000 watt servers @ 80% Load) | Energy cost (S.08 kW/hr) |
|---|---|---|
| 77 | 1,000 x (1,000 x .8 load) /.77 = 1,048 kWh | 83.84 |
| 88 | 1,000 x (1,000 x .8 load) /.88 = 909 kWh | 72.72 |

Increasing power supply efficiency from 77 to 88% generates savings of $11 an hour or $97,000 annually.

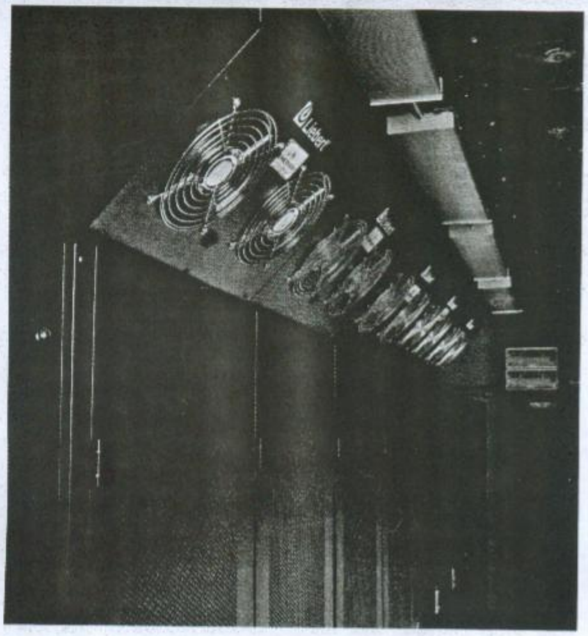

Manufacturers have been developing creative and competitive approaches to cooling data centers.

be available. The local utility or PUC may have experts available to help develop a demand reduction program.

## Power Management

Focusing on the IT side of the equation, the first and most important question to ask is, "What is the utilization of IT equipment and how does it vary based on time of day or season?"

Jeff Bezos, CEO of Amazon, recently revealed that his company's IT equipment utilization was just 10% during some periods of the year. Average data center utilization is typically around 15 to 20% when calculated on a 7 by 24 by 365 basis. Can energy efficiency be optimized during off-peak periods without impacting performance during peak periods?

The answer is a resounding yes. Start with pending IT equipment purchases. Prior to purchase is the best time to drive for efficiency in equipment specification. Work with suppliers to understand the efficiency of new equipment and how it compares to legacy systems.

Efficiency at the server level is a moving target depending upon supplier.

The 80plus.org's specification that recommends that all server power supplies have 80% or better efficiency from 20% of load up. There is also a pending EPA Energy Star server test protocol designed to characterize a server's efficiency over a no-load (zero utilization) to full-load (100% utilization) range. Note that at zero utilization, many of today's servers still draw close to 85% of full-load power so improvements in power supply efficiency generate substantial savings, even when a data center is lightly loaded.

Changing servers from a 77% efficient power supply to one nominally rated at 88% (90 to 92% efficiency power supplies are available at power levels greater than 3 kW), generates savings that more than pay for the cost of the improved power supply within the first year, assuming a cost runs $100 to $150 per power supply (see the table).

All new product specifications should include an operating efficiency specification. This should reflect a wide-range of use, not a single optimal test point. Though not yet released, the EPA server-test protocol is a good reference document to incorporate within your IT equipment specification. Again, utilities and PUCs may extend technical resources to assist in demand-side reductions. Sun Microsystems and PG&E recently announced the industry's first rebates for energy-efficient IT equipment.

Legacy equipment can also benefit from these advancements. On many servers, the process of swapping out lower-efficiency power supplies for new higher-efficiency supplies is relatively simple and can be handled by an IT technician in minutes. Or, a business can work with its supplier(s) on a platform consolidation and upgrade. This may get a bit pricey but allows a worker to never find out that you can collapse 500 legacy servers onto 100 new servers with a new operating system and other enhancements.

But, what about all those servers consuming 80% of full-load energy while sitting idle through the night or on weekends? With a data center running 24 by 7 and a business that runs 12 by 5½, not counting holidays, there are huge opportunities for energy reduction. A CEC/LBNL study found data centers tend to run nearly flat-line in power consumption regardless of time of day or day of week, etc. This indicates that IT professionals are not taking advantage of the latest power management features available from AMD, IBM, Intel, or Sun that enable processors to change P-states depending upon utilization.

The reason is probably fear that the processors may not ramp back up to full speed and may miss a significant transaction. That may be reasonable during normal business hours, but there is no good reason to run 100% of servers flat-out during predictably slow periods of the day or week.

Assuming a worst-case model of wanting to leave 50% of servers idling in full-speed,

Data centers are a power-hungry necessity of today's business world, but they are rapidly coming under the scrutiny of several important groups.

full-power mode to ensure nothing is missed in the off-hours, this still leaves 50% that can be throttled-back to a lower power state. These servers are not powered off; rather, they are taking advantage of new power management capabilities. Server and chip OEMs will gladly demonstrate the reduced state processing power and ramp-up characteristics that demonstrate a quick recovery time.

In this throttled-back mode, servers can achieve a 30 to 50% energy savings depending upon memory, storage, and other I/O characteristics. Given a thousand or more 500 to 1,000-W servers, this represents a huge potential savings. For example, a typical 7 a.m. to 7 p.m., 5½-day week business has about 103 hours a week where the bulk of servers can be brought into an energy managed mode. Looking at 500 servers (50% of 1,000 nominal units) at 1,000 watts and energy costs of $.08 kW/hr, savings of $10.90 an hour are achieved. This adds up to $1,123 a week or approximately $55,000 a year, with absolutely no reduction in performance or increase in risk.

## Optimal Asset Location

Lastly, location is important. No the data center can't be moved, but it can moved in the dataline. This assumes someone is already implemented the hot-aisle/cold-aisle rack arrangement and is looking for even greater operational efficiency.

Examine the current room layout from a density perspective. Where are the highest power/thermal density assets? Depending upon room constraints, there are a few potential scenarios that will enable system operators to more closely couple cooling and power delivery and gain significant operating efficiencies. When build out plans call for ever higher power and cooling densities, it may become time to begin to deploy a hybrid cooling architecture that introduces localized, supplemental cooling at the rack (liquid or refrigerant-cooled rack enclosures or overhead systems) or row level (aggregated liquid cooled racks or over rack /row systems). These supplemental cooling systems run in conjunction with base level HVAC/precision cooling systems but at a far greater level of cooling per watt.

Such a cooling scheme may yield cooling energy savings of 20 to 30% over traditional installed infrastructure, thereby greatly improving PUE/DCE ratio while further reducing operating expense.

## Conclusion

Data centers are a power-hungry necessity of today's business world, but they are rapidly coming under the scrutiny of several important groups. Government and public agencies along with utilities and environmental concerns are beginning to investigate operational practices within the data center from the core server/storage device all the way out to the campus-level infrastructure and operations. End users and OEMs can take steps today to optimize equipment and facility energy efficiency to stem both rising operating costs and increased regulation.

**About the Author: Jack Pouchet**, director Marquee Accounts for Emerson Network Power/Liebert, works closely with major OEMs in the server market, large data center users in the financial and communications industries, and leading mission critical engineering firms to help identify new trends and opportunities for advanced power and cooling technologies. ***e&pm***